Energy efficiency in access and aggregation networks: From current traffic to potential savings

Original

Energy efficiency in access and aggregation networks: From current traffic to potential savings / Edoardo, Bonetto; Finamore, Alessandro; Mellia, Marco; Riccardo, Fiandra. - In: COMPUTER NETWORKS. - ISSN 1389-1286. - STAMPA. - 65:(2014), pp. 151-166. [10.1016/j.comnet.2014.03.008]

Availability:

This version is available at: 11583/2543102 since:

Publisher:

Elsevier

Published

DOI:10.1016/j.comnet.2014.03.008

Terms of use:

openAccess

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Energy Efficiency in Access and Aggregation Networks: from Current Traffic to Potential Savings

Edoardo Bonetto^{a,*}, Alessandro Finamore^a, Marco Mellia^a, Riccardo Fiandra^b

^aPolitecnico di Torino
Corso Duca degli Abruzzi 24, 10129 Torino - IT
^bFASTWEB
Via Caracciolo 51, 20155 Milano - IT

Abstract

Access and aggregation networks account nowadays for a large share of the consumed energy in communication networks, and actions to ameliorate their energy cost are under investigation by the research community. In this work, we present a study of the possible savings that could be achieved if such technologies were in place. We take advantage of large datasets of measurements collected from the network of FASTWEB, a national-wide Internet Service Provider in Italy. We first perform a detailed characterization of the energy consumption of Points of Presence (PoPs) investigating on how factors such as external temperature, cooling technology and traffic load influence the consumed energy. Our measurements precisely quantify how the power consumption in today networks is practically independent from the traffic volume, while it is correlated only with the external temperature. We then narrow down our analysis to consider the traffic generated by each household. More specifically, by observing about 10,000 ADSL customers, we characterize the typical traffic patterns generated by users who access the Internet.

Using the available real data, we thus investigate if the energy consumption can be significantly reduced by applying simple energy-efficient policies that are currently under studies. We investigate energy-to-traffic proportional and resource consolidation technologies for the PoP, while sleep modes policies are considered at the ADSL lines. All these energy-efficient policies, even if they are not yet available, are currently being widely investigated by both manufacturers and researchers. At the PoP level, our dataset shows that it would be possible to save up to 50% of energy, and that even simple mechanisms would easily allow to save 30% of energy. Considering the ADSL lines, it results that sleep mode policies can be effectively implemented, reducing the energy consumption of ADSL modems with little or marginal impact on the Quality of Service offered to users. We make available all datasets used in this paper to allow other researchers to benchmark their proposals considering actual traffic traces.

Keywords: energy-efficiency, access networks, sleep mode

1. Introduction

The topic of energy efficiency in telecommunication networks has attracted a lot of interest in the recent years. This is due to economical reasons, driven by the ever increasing cost of energy, to the willingness of reducing the environmental impact, and to technical matters, since energy density limits devices scalability.

In the telecommunications field, researchers were used to consider energy-efficiency as an aside aspect of their investigations. Instead, nowadays several initiatives such as TREND¹ and ECONET² European projects, the GreenTouch Consor-

tium³ and the D-Link Green Products⁴, focus specifically on the energy efficiency of telecommunication devices. Among those initiatives, the TREND project aims at quantitatively assessing the energy demand of current and future telecommunication infrastructures, and designing energy-efficient, scalable and sustainable future networks. In this work we present the results of a collaboration between Politecnico di Torino and FASTWEB⁵, one of the major Internet Service Provider (ISP) in Italy.

We focus explicitly on wired networks. In such scenario, the access and the aggregation portions of the network account for the largest share of the overall energy consumption, and forecasts show that this is expected to worsen in the future [1]. The improvement of the energy efficiency in these network segments is as such an important target.

In this work, we aim at (i) providing a detailed characterization of the actual energy consumption as seen by the ISP, and (ii) evaluating the impact of currently studied energy efficient

*Corresponding Author: Edoardo Bonetto, Department of Electronics and Telecommunications, Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129 - Torino, Italy email: edoardo.bonetto@polito.it, tel: +39-011-090-4229, Fax: +39-011-090-4099

Email addresses: edoardo.bonetto@polito.it (Edoardo Bonetto),
alessandro.finamore@polito.it (Alessandro Finamore),
marco.mellia@polito.it (Marco Mellia),
riccardo.fiandra@fastweb.it (Riccardo Fiandra)

¹<http://www.fp7-trend.eu>

²<http://www.econet-project.eu>

³<http://www.greentouch.org>

⁴<http://www.dlinkgreen.com/greenproducts.asp>

⁵<http://company.fastweb.it>

policies when effectively deployed in these network segments. Our goal is to provide a precise quantification of the possible energy saving that policies and technologies currently under development would bring in a real scenario. For this purpose, we leverage a large dataset containing the actual energy consumption, traffic and other important measurements collected by the operative network of FASTWEB. *Using these actual data, we analyze how the energy consumption is correlated to the other metrics in order to understand how and where energy efficiency can be improved. The characterization of the energy relationships with other factors represents indeed an important aspect of this work. Even if these relationships are intuitive, it is fundamental to actually confirm and quantify it with actual data.* Following the scientific principle of allowing results reproduction, and to allow other researchers to validating their proposals in actual scenarios, we make available all datasets used in this paper on the website of the TREND project ⁶.

In the context of access and aggregation networks, we focus our attention on the Points of Presence (PoPs), large network “nodes” that act as both aggregation and traffic switching points, where several thousands of end-users lines are aggregated. This choice is motivated by the fact that the energy consumption of PoPs weights for the 26% of the total energy consumed by FASTWEB, as we computed starting from the data provided by FASTWEB. We next move our attention to each single ADSL access line, i.e., on the modems located at users’ home and at the ISP Digital Subscriber Line Access Multiplexer (DSLAM). Being in the order of several millions, it has been indeed shown that the customer premise equipment overall accounts for more than 65% of the energy consumed in access networks [2].

Notice that the PoP and the single-line operate in substantial different traffic scenarios and time scale. The former transports traffic aggregated from thousands of users, while the latter manages traffic generated by a single household. This have to be taken into account when conceiving energy efficient policies. Having this in mind, we investigate some “what-if” scenarios in which we identify possible energy efficient policies that could be deployed if new technologies become available, or if different management procedures are adopted.

We assume the availability of energy efficient policies, based on resource consolidation or sleep mode technologies that are currently being investigated by manufacturers and researchers. We explicitly do not focus on the problems related to their technical implementation and deployment. We rather evaluate if these strategies could offer significant energy savings assuming as benchmark the available measurements. In the case that savings are consistent, manufacturers can be motivated to perform the investment to provide those solutions, and ISPs to deploy them.

By using simple models that we fed by the actual data, we predict the possible energy savings that can be expected when rolling out such technologies. Results show that even simple policies allows to save from 30% to 50% of the energy consumed by a PoP. Furthermore, given the very large burstiness

of traffic on each single ADSL line, sleep mode policies would allow to keep the ADSL modes into sleep or low-power state for more than 60-80% of the time with a marginal impact on users’ perceived Quality of Service (QoS). Technologies such as the one proposed in [3] would be an excellent means to reduce the energy costs of access networks.

Preliminary results have been presented in [4, 5]; in this paper we extend them introducing a more comprehensive analysis of the PoP traffic characteristics and considering different energy efficiency technologies. *In particular, we perform a more detailed analysis of the dataset, introducing the analysis of Pearson correlation coefficients in order to better characterize the correlation among the different metrics. Moreover, we perform a detailed analysis of the energy breakdown of the PoPs in order to distinguish between networking devices and cooling system power consumption contributions. Furthermore, we consider the adoption of Dynamic Voltage Frequency Scaling (DVFS) based technologies as PoP strategy to improve the energy efficiency. Achievable savings in term of energy consumption and monetary investments are also quantified for the selected PoPs and the associated ADSL lines.*

In the following, Sec. 2 describes the dataset and the available metrics. In Sec. 3 we present the energy profiling of the PoPs, evaluating possible correlations among energy consumption and the other metrics, and furthermore we introduce the energy breakdown of the PoPs. We analyze the traffic dynamics of the PoPs and we evaluate the typical activity of ADSL user in Sec. 4. The energy saving strategies are introduced in Sec. 5, while the achievable energy savings are estimated in Sec. 6. Lastly, conclusions are drawn in Sec. 7.

2. Dataset

In this paper we leverage actual datasets collected from the operative network of FASTWEB. Considering PoPs, we focus on an extensive set of informations including (i) their characteristics, (ii) their energy consumption, (iii) the outside temperature and (iv) the total traffic they carry. Considering single ADSL lines, we characterize the traffic generated by about ten thousand end customers. In the following, we detail the measurements that are instrumental for our analysis.

Each PoP is characterized by the number of connected customers, the number of hosted networking devices, e.g., Routers, DSLAMs, B-RAS (Broadband Remote Access Server), switches, etc., the type of technology, i.e., ADSL or FTTH, and the type of air conditioning system. Since FASTWEB has several PoPs, we have chosen those that present different characteristics considering size, location, access technologies and air conditioning systems.

The energy consumption and the outside temperature measurements have been collected thanks to a management system that FASTWEB has recently deployed. The goal of this system is to monitor the energy consumption of the company facilities to highlight energy consumption changes due to failures or to misconfigurations. In details, the system provides:

- **Total energy consumed by a PoP:** Let $E_P(i)$ be the energy consumed by PoP P during the i -th time window. It comprises

⁶<http://www.fp7-trend.eu/content/datasets-benchmark-access-networks>

the energy consumed by all devices, including network equipment and air conditioning systems. The energy is expressed in kWh and a measure is available every $\Delta = 15$ min. To avoid outliers due to possible misreading, we restrict the set of measurements to values comprised within the 1-st and 99-th percentile. We denote this set as $\mathcal{E}_P$.

- **External air temperature:** Let $T_P(i)$ be the external temperature in Celsius degree during the i -th time window for the PoP P . A sample is available every $\Delta = 1$ hour. We denote this set as $\mathcal{T}_P$.

- **Normalized bit rate:** Let $B_P(i)$ be the total amount of traffic processed by PoP P in $\Delta = 15$ min time interval. It corresponds to the sum of all the data exchanged between the customers and the network. Thus, the average bit rate in time window i is given by $B_P(i)/\Delta$. We report the normalized value assuming that the maximum is equal to 0.8. This assumption has been made in order that each PoP is not operating at its maximum capacity, but it is little over-dimensioned as it is usual done in reality. This results in,

$$\rho_P(i) = 0.8 \frac{B_P(i)/\Delta}{\max_i(B_P(i)/\Delta)} . \quad (1)$$

We denote this set as $\mathcal{R}_P$.

Concerning the measurements of the traffic processed by the PoP, we consider the period starting from January 2011 until the end of March 2011. We restrict our analysis on traffic to this period since the number of users at each PoP can be considered stable, while it typically changes during other periods of the year when customers are attracted by special offers. Instead, we collected measurements of energy and external temperature in the period between mid September 2010 and mid September 2011 in order to cover all possible seasons. Since the temperature samples are available with a time interval Δ of 1 hour, we assume that the temperature is not changing between measurements and, in case of no sample available, we consider the last measured one.

The traffic measurements of each single ADSL line consist in:

- **Number of exchanged bytes:** Let $bytes_U(i)$ be the total amount of bytes sent and received by the user U in the i -th time interval of $\Delta = 10$ seconds. These samples have been collected using the traffic monitoring tool Tstat⁷ [6]. We compute the average line rate, expressed in b/s, as $bytes \cdot 8/\Delta$.

We collect traffic data of about 10,000 ADSL residential customers associated to the same PoP, for 24 hours, starting from 31 May 2011 at 11pm.

3. Energy consumption characterization of the Points of Presence

In this section we characterize the energy consumption of PoPs, evaluating the correlation with system characteristics, traffic, and external temperature. We then present the energy consumption breakdown of each PoP, in order to estimate the

ID	Users [k]	Devices	FTTH [%]	Free cooling [%]	$r^{(temp)}$	$r^{(load)}$
A	3.8	6	100	100	0.9	0.24
B	52.2	31	33	0	0.47	0.11
C	22.8	35	0.2	12	0.81	0.09
D	16.3	19	86	0	0.8	0.08
E	13.5	34	92	0	0.75	0.28
F	62.3	72	11	40	0.75	0.35
G	61.9	44	6	100	0.85	0.36
H	22.5	32	0.2	4	0.75	0.13

Table 1: Summary of PoPs characteristics

energy efficiency of the networking devices and of the air conditioning system.

3.1. Overview of the main characteristics of the PoPs

The main characteristics of the PoPs that we selected for our study are summarized in Table 1. Hereafter, these PoPs are identified with a capital letter ranging from A to H. PoPs have different size (number of users and networking devices they host) depending mainly on the area that they cover. The number of users ranges from a few thousands for PoP A to more than 60 thousands for F or G. Similarly, the number of devices ranges from a few units of PoP A to more than 70 of PoP F. **The networking devices present in a PoP consist mostly in traffic aggregation devices which aggregate the users traffic received from DSLAMs, or from the FTTH access networks. They provide access to the operative network of FASTWEB by means of routers and B-RAS. In PoP E and F some backbone devices (backbone routers, WDM multiplexers) are also present.**

FASTWEB offers end-customers both ADSL and FTTH as access technology. Depending on the infrastructure that has been deployed in the surroundings of the PoPs, both or just one between ADSL or FTTH systems can be present in a PoP. Column four reports the percentage of users that uses FTTH access technology. Note that typically the PoPs have a predominant access technology.

Inside each PoPs, air conditioning systems are present. They belong to two different technologies: heat pump, based on compressors only systems, and free cooling supported systems. In the latter case, when the external temperature is sufficiently low, there is the possibility to cool the inside of the PoP using the outside air without the help of energy-hungry heat-pump solutions. The percentage of free cooling capacity of a PoP are detailed in the fifth column of Table 1.

3.2. Energy consumption correlation

The last two columns of Table 1 report the Pearson correlation coefficients [7] of the energy consumed by a PoP in relation, respectively, with the external temperature and with the processed traffic.

We rely on the sample Pearson correlation coefficients, denoted by r in Eq. 2. Given two datasets $\mathcal{X}$ and $\mathcal{Y}$, each of size N , let $x_i \in \mathcal{X}$ and $y_i \in \mathcal{Y}$ be the samples collected at the same time interval i . Let $\bar{x}$ and $\bar{y}$ be the average value computed on

⁷<http://tstat.polito.it>

Figure 1: Energy consumption and external temperature versus time.

Figure 2: Energy-temperature correlation coefficient vs. free cooling capability.

all samples in $\mathcal{X}$ and $\mathcal{Y}$ respectively. The sample Pearson correlation coefficients r is computed as:

$$r = \frac{\sum_{i=1}^N (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^N (x_i - \bar{x})^2} \cdot \sqrt{\sum_{i=1}^N (y_i - \bar{y})^2}}. \quad (2)$$

We have that $r \in [-1 : 1]$. Values close to 0 means that the two metrics have none or little correlation. Values close to 1 (-1) are positively (negatively) correlated.

3.2.1. Correlation with external temperature

For a PoP P , $r_P^{(temp)}$ is computed using $\mathcal{X} = \mathcal{E}_P$ and $\mathcal{Y} = \mathcal{T}_P$. Notice the strong correlation for most of the PoPs. In particular, the correlation is stronger in the PoPs where the percentage of air conditioners with the free cooling is larger. To qualitatively gauge the correlation, we take PoP A and PoP B as example; those are the PoPs with the highest and the lowest correlation coefficient respectively, and they are the PoP with complete and no free cooling capability. Figs. 1(a) and 1(b) report the energy and temperature of the whole measurement campaign for both PoPs. Energy consumption and temperature are reported in the top and bottom part of each plot, respectively. Looking at Fig. 1(a), it can be noticed that energy and temperature follow the same seasonal behavior, with low values from October to March⁸. In particular, during the coolest days of the year, the energy reaches its minimum values; when spikes in the temperature are present, the energy consumed is larger. On the contrary, Fig. 1(b) shows that, for PoP B, the consumed energy barely follows the yearly temperature trend.

We further investigate the impact of the air conditioning system over the correlation in order to give a qualitative idea of how much this coefficient depends on the free cooling. The relationship is shown in Fig. 2, which depicts the correlation coefficient $r^{(temp)}$ of all PoPs versus the percentage of free cooling air conditioners capacity. In summary, energy consumption is highly correlated with external temperature; however, even if a little trend seems to be present, it is not possible to state that there is a clear relationship between the energy-temperature correlation coefficient and the percentage of free cooling. Indeed, several other factors influence this correlation, such as the exposure and the physical location of the PoP, how the building hosting the PoP is built, the number and type of hosted devices, etc. Unfortunately, the collected data do not allow to study such variables. For instance, PoP F presents a similar correlation coefficient to other PoPs even if it has a higher percentage of free cooling. A possible reason could be that PoP F is the largest PoP, considering number of users and of networking devices, or this difference can be due to the fact that the PoP contains also some backbone devices, while other PoPs, apart PoP E, do not.

3.2.2. Correlation with traffic load

The correlation $r_P^{(load)}$ between the energy consumed and the traffic load is computed, for a PoP P , using $\mathcal{X} = \mathcal{E}_P$ and $\mathcal{Y} = \mathcal{R}_P$. It is reported in the last column of Table 1.

$r_P^{(load)}$ values are much smaller than $r_P^{(temp)}$, meaning that there is smaller correlation between energy and traffic. To qualitatively appreciate this, we report the daily variation of consumed energy and traffic load of PoP B in Fig. 3. Two days are selected; the first day refers to the winter (black lines), and the second day to the summer (gray lines). Left y-axis reports the normalized traffic (dashed lines), while right y-axis reports the energy consumption (solid lines). While the traffic profile fol-

⁸The sudden jump at beginning of April in Fig. 1(a) is due to a change in the configuration of the cooling system.

Figure 3: Energy consumption and normalized traffic profiles of PoP *B* during one day.

lows the traditional day-night pattern (see Sec. 4 for more detailed analysis), the energy consumption profiles present a completely different and uncorrelated behavior with respect to the traffic.

Furthermore, some interesting observations about the correlation of energy and external temperature can be deduced. The energy consumption presents an intra-day periodic behavior that can be explained by the switching on and off of the air conditioner compressors according to a well defined scheme. In the warmer days, an additional effort is required to cool down the PoP. Notice indeed the smaller number of peaks since the conditioners are switched off less frequently. During the colder season, conditioning system consumes less energy as reflected by the more frequent peaks corresponding to air conditioners becoming idle.

3.2.3. Correlation with other indexes

We have also studied the correlation between the energy consumption of the PoPs and other parameters. In particular, we have considered (i) the size of the PoP in terms of number of supported users, and (ii) the predominance of ADSL or FTTH technology.

Fig. 4 shows the per user energy consumption. Measurements do not allow to conclude that there is correlation between energy consumption and size of the PoP. For example, in both PoP *A* and *B*, the average consumption is about 0.5 Wh/user despite PoP *A* accounts for only one tenth of the users served by *B*. Notice that the high values of PoP *E* and *F* are due to the presence in those PoPs of some backbone devices (backbone routers, WDM multiplexers) that are not present in other PoPs. Similarly, no trend can be associated to the access technology. Results are not reported here for the sake of brevity.

3.3. Energy breakdown analysis

We are interested to evaluate the impact of air condition energy cost with respect to the energy consumption of networking devices. It is important to understand how consistent is the part of energy that can be attributed to the networking equipment. Indeed, it may be convenient to deploy new energy-aware

Figure 4: Energy consumption per user versus number of users.

ID	min($\mathcal{E}$) [kWh]	max($\mathcal{E}$) [kWh]	<i>PUE</i>
<i>A</i>	0.7	3.7	5.29
<i>B</i>	14.28	16.75	1.17
<i>C</i>	10.34	14.56	1.41
<i>D</i>	6.18	7.53	1.22
<i>E</i>	13.63	16.84	1.24
<i>F</i>	41.4	59.42	1.43
<i>G</i>	21.44	29.13	1.36
<i>H</i>	12.41	17	1.37

Table 2: Energy breakdown of the PoPs

strategies, involving networking devices, in case those are responsible for a large fraction of energy costs.

Unfortunately, the energy management system reports only aggregated measurements, i.e., it does not separate network devices power consumption from the one of air conditioning systems. To obtain this breakdown, we assume that the minimum energy consumption measured on the whole dataset represents a good estimation of the power consumed by networking devices. **If air conditioners are not working, the measured energy consumption can be assumed to correspond to the energy consumed by the networking devices only.** Indeed, since the air conditioners automatically are turned off when the temperature goes below the desired target, it is expected that at given time instants there will be no conditioners working. In those time instants the energy consumption measured will have the minimum value since only the energy consumption of networking devices will be present. **This energy measure can be taken as estimation of the consumption of networking devices. We verify this assumption later.**

We evaluate the energy efficiency of the PoPs considering the Power Usage Effectiveness (*PUE*) [8, 9]. This metric is computed as the ratio between the power consumed by the overall facility and the power consumed just by the IT equipment. A value of $PUE = 1$ indicates that all energy is consumed by IT devices. $PUE = 2$ means that for every watt of IT power, an additional watt is consumed to cool the networking equipment. While several measurements of *PUE* are available for data centers [10], to the best of our knowledge, no measurements have been presented considering networking facilities. In our sce-

nario, PUE is estimated by the fraction of the maximum and the minimum value of energy consumption for each given PoP dataset, i.e., for the PoP P , have

$$PUE_P = \frac{\max(\mathcal{E}_P)}{\min(\mathcal{E}_P)}. \quad (3)$$

Recall that $\mathcal{E}_P$ is the set of samples within the 1-st and the 99th percentile of the $E_P(n)$ distribution to avoid outliers.

FASTWEB, due to logistic constrains, instrumented only PoP B to cross check the PUE estimation. The difference among the measured and the estimated PUE is of only around the 13%. Thus, measurements show that $\min(\mathcal{E}_B)$ is a good estimate of the energy consumption of the networking devices only.

Table 2 reports the estimated PUE for all PoPs. Notice that the PUE ranges from about 1.2 to 1.4, i.e., air conditioning costs less than 30% of PoP energy. As a reference, Google Inc. claims to have the most efficient data centers with a PUE of 1.12⁹. Also the data centers of Facebook are achieving similar PUE¹⁰.

PoP A is an outlier with $PUE = 5.29$. This is due to the fact that PoP A is a new PoP equipped with just six networking devices (see Table 1) while the air conditioning system has been (over)dimensioned to scale for future installation.

Main findings

The energy profiling of the PoPs leads to the conclusion that their energy consumption is strongly correlated to the external temperature. This correlation is due to the air condition system which accounts for less than 30% of the PoP total energy consumption.

Although traffic is highly variable during the day, the energy consumption exhibits no correlation to the processed traffic. This is in line with the measurements of single devices [11] that exhibit marginal proportionality between energy consumption and traffic load.

4. Traffic characterization of Points of Presence and ADSL lines

Energy efficient policies based on energy-to-traffic proportionality can be effective if the traffic volume exhibits periodicity over time. At the moment, energy-to-traffic proportionality is not offered by devices and, as our findings confirm, the energy consumption is not correlated to the traffic volume. In this work, we want thus to precisely quantify how much could be the energy saving if energy-to-traffic proportionality will be introduced at the PoPs.

In this section, we focus on the traffic processed by PoPs in order to verify if the traffic varies sufficiently over time. Intuitively, a PoP is aggregating the traffic of a large number of

Figure 5: Typical daily pattern at a PoP along one week.

Figure 6: Daily traffic distribution for a week of PoP A .

users; thus it is expected that the traffic follows a daily profile without having abrupt variations. When instead an end-customer access line is considered, the traffic depends on the activity of the users. We expect that traffic is much more bursty, and the profile can vary significantly from a line to another one.

The precise quantification of the above intuitions is the goal of this section. We consider this as a major contribution of this paper, and as such, we offer the dataset to the research community to allow other researchers to benchmark their proposals.

4.1. Macroscopic traffic analysis of the PoPs

Let us start by focusing on the traffic variation of the PoPs during a one-week long period. Fig. 5 reports measurements for PoP B , being the patterns for the other PoPs very similar. The day/night periodicity is clearly visible, with off-peak traffic that is about one third of peak traffic.

To characterize the minor differences among different days, we proceed as follows. Let $\mathcal{R}_P(i)$ be the set of PoP P traffic measurements that correspond to day i . Fig. 6 shows the marginal Cumulative Distribution Function (CDF) of $\mathcal{R}_A(i)$, which is representative of other PoPs. Measurements cover the 7 days of a week; Saturday and Sunday are highlighted. As it can be clearly noticed, week-days are practically all equal, while Saturday and Sunday show that the PoP carries a larger fraction of traffic for longer time period. This reflects the end-users' weekly habits.

⁹<http://www.google.com/about/datacenters/efficiency/internal/>

¹⁰<https://www.facebook.com/ForestCityDataCenter/>,
<https://www.facebook.com/PrinevilleDataCenter/>

Figure 7: Traffic distribution of every PoP during 3 months.

Figure 8: Examples of the evolution of the line rate for three different users.

We repeat the measurement by considering the whole three months period and for all PoPs. Fig. 7 depicts the resulting profile. It confirms that traffic follows the same behavior for all PoPs, i.e., the aggregate behavior of offered traffic follows the same daily pattern in a quite stable and easily predictable way, regardless the location, size and kind of PoP.

Finally, we measure the variability of traffic as the ratio between the peak and off-peak traffic volume. To consider a conservative case, we compute, for each day and for each PoP, the ratio among the 80-th and 20-th percentile of the traffic measurements of the day. **This ratio indicates how much the traffic volume varies at a PoP during the day. We found that all PoPs have a ratio greater than 3, meaning that the traffic volume at the peak time is more than 3 times the traffic volume at the off-peak time. This result indicates that there is a significant difference between the peak and the off-peak traffic volume which is an important condition for the effectiveness of the energy-to-traffic proportional based policies.**

4.2. Microscopic traffic analysis of ADSL lines

We now consider measurements at the granularity of single customer access line. Typical examples of users' activity are reported in Fig. 8, which reports the line rate for a 24h period. Notice the logarithmic scale on the y-axis. Three different usage patterns are exposed: *user1* is always connected constantly exchanging a large volume of traffic (as typical of users running

Figure 9: CDF of the total activity of the monitored lines during the day.

P2P applications); *user2* has two activity periods in the early morning and afternoon during which the bit rate is very large; few activity periods are possibly due to the usage of Voice over IP (VoIP) services; finally, *user3* shows only three significant traffic events during which the bit rate is limited. These are possibly due to VoIP as well.

Interestingly, both *user2* and *user3* exhibit several traffic events, in an almost periodic way, characterized by very low bit rate. Those correspond to traffic automatically generated by applications for control purposes, e.g., the DHCP traffic due to the modem itself, or applications that check for the availability of updates. In the following, we refer to this type of traffic activity as *background traffic*. Note that during time periods in which only background traffic is present, no user is actually actively accessing the network. Thus the line could enter into sleep mode or low power/low rate states in order to save energy.

The deployment of sleep mode policies in the access lines can be a good solution to improve the energy efficiency in access networks, but their effectiveness depends on which is the typical user's behavior. Indeed, if the majority of users has habits similar to *user1*, sleep mode policies would result useless. In the case that the users' habits are similar to *user2* and *user3*, there is room to achieve significant energy savings.

In the following, we characterize the typical end-customer behavior by exploiting our measurements. We define as "activity period" the time during which the line rate is above a certain threshold *MinRate*. We consider that all the traffic events for which the line rate is below *MinRate* are not due to user's activity, but they represent background traffic.

Considering 24h time interval, Fig. 9 reports the CDF of the total line activity period for different values of *MinRate*. Overall, the average line utilization is very limited. For instance, if no threshold is applied, only 50% of the lines is active for more than 2.5 hours, and only 10% of lines exchanges traffic for more than 12 hours. Considering *MinRate*=1 kb/s, 50% of the lines is active for less than 1 hour in the entire day. Even more impressive, we have that about 30% of the lines never exceeds *MinRate*=100 kb/s in the whole day.

Main findings

The analysis of actual traffic traces shows that there are two main areas where it is worth to apply energy saving policies. At the PoP level, the aggregated traffic profile presents high

dynamics; a regular day-night pattern is followed, with the off-peak load that is more than three times smaller than peak load. At the single ADSL line, most of customers actually **generate** traffic for typically short periods, with their access line that carries only background traffic for most of the time.

5. Energy efficient policies, a what-if analysis

Given the traffic characteristics of PoPs and single end-users' line, different approaches can be envisaged to deploy energy-wise policies. In particular, PoPs are better suited for the deployment of energy proportional technologies since connectivity must be guaranteed at any time. Single lines are instead more suited for sleep mode policies, according to which the DSLAM and home modems enter sleep mode, possibly losing connectivity.

The aim of this section is to evaluate which are the benefits that these strategies may bring.

5.1. Energy proportionality in the PoPs

The energy proportionality refers to the capability of a device, or a set of devices, to adapt the consumed energy to the actual traffic load. We consider two different approaches to tackle the issue of improving the energy efficiency: (i) the availability of a new generation of networking devices which offer *energy-to-traffic proportional technologies*, (ii) a modular organization of the devices in the PoPs according to the *resource consolidation* practice [12, 13, 14].

In the first case, it is necessary to wait for technological advances. As seen in Sec. 3.2.2, current devices present an energy consumption that is almost independent from the traffic they handle.

In the second case, instead, the organization of networking devices in a PoP should be revised. For example, during low traffic periods, a set of devices could be switched off, leaving on only a minimal set of devices able to satisfy the current traffic requirements.

Both strategies are not yet deployed in current networks. In this work, we want thus to understand if these strategies can effectively help in reducing the energy consumption. We thus perform a “what-if” analysis considering the future deployment of these strategies.

We account only networking devices energy consumption optimization, neglecting air conditioning systems. However, since the emitted heat depends on the energy consumption of the devices, it can be expected that some savings can be achieved also by the conditioning system.

5.1.1. Energy to traffic proportional technologies

We consider the availability of networking devices having an energy consumption which is fully or partially proportional to the traffic load. Let $E(\rho)$ be the normalized energy consumption for a given offered load ρ . Following a classical model, we assume $E(\rho)$ is the sum of a static part, E_{static} , independent from

the traffic, and of a dynamic part, $E_{dynamic}(\rho)$, function of the traffic load ρ . That is

$$E(\rho) = E_{static} + E_{dynamic}(\rho). \quad (4)$$

In particular, we consider that the dynamic component of the energy is *linearly* or *cubicly* proportional to the traffic processed. The latter represents an approximation of implementing DVFS technologies in the devices, since the consumption scales cubically with the operating frequency [15].

In the following, we assume to use normalized values for the traffic load ρ and for the energy consumption values. The normalized traffic of the PoPs is computed according to Eq. 1. Moreover, we assume that the normalized energy consumption of the current technology (no proportionality) is equal to 1, and that is also the consumption at maximum traffic load for any energy proportional strategy, thus $E(1) = 1$.

Thus, in the case of *linear* proportional energy consumption, the actual consumed energy can be computed as:

$$E(\rho) = E_{static} + \alpha \cdot \rho \quad (5)$$

where $\alpha \cdot \rho$ represents the variable component of consumed energy, $\alpha \in [0 : 1]$ being the proportionality factor.

Since we consider the normalized energy $E(1) = 1$, the fixed energy cost becomes $E_{static} = 1 - \alpha$. The case $\alpha = 0$ corresponds to current technology, whose consumption is constant, regardless the load, and thus the normalized consumption is 1.

Similarly, when $E_{dynamic}$ is a cubic function of the traffic load, we have that:

$$E(\rho) = E_{static} + \alpha \cdot \rho^3. \quad (6)$$

Imposing $E(1) = 1$, we have also in this case $E_{static} = 1 - \alpha$.

Finally, according to the previous assumptions, the total energy savings S that can be achieved, is computed as:

$$S = 1 - \frac{1}{N} \sum_n E(\rho(n)) \quad (7)$$

where N is the total number of samples.

5.1.2. Resource consolidation

In this approach, the networking devices in a PoP are deployed with a modular organization and, during low traffic demand periods, a fraction of them can be put into sleeping mode.

For instance, consider a PoP with two possible operating states corresponding to High and Low capacity, with High/Low states entered when traffic goes above/below a given threshold. This strategy can be generalized for more than two states. In this work, we consider in details the cases for which PoPs can operate in two or three different operating states. Fig. 10 shows examples of the different energy saving schemes.

In the following, for simplicity, we assume that when the traffic is above the fraction m of the peak traffic, the PoP is fully operative and consumes the maximum energy, that is normalized to 1. When traffic is below the fraction m of the peak, the PoP can work at a fraction m of its full capacity, and it is also consuming a fraction m of the consumption at full capacity. We

Figure 10: Energy savings schemes: no proportionality, full load proportionality, 2-states and 3-states, from left to right.

refer to this case as *two-states* resource consolidation. We have that:

$$E(\rho) = \begin{cases} m & \rho < m \\ 1 & \rho \geq m \end{cases} \quad (8)$$

Finally, we extend the previous case and assume that a PoP works in *three-states*: one for low traffic, one for medium and one for maximum traffic. Thus,

$$E(\rho) = \begin{cases} l & \rho < l \\ m & l \leq \rho < m \\ 1 & \rho \geq m \end{cases} \quad (9)$$

where l and m are the traffic thresholds and corresponding energy consumption values for low and medium traffic, respectively.

5.2. Sleep mode policies at the ADSL lines

The analysis reported in Sec. 4 shows that ADSL users are typically inactive for long time periods. Thus it would be advantageous to implement sleep mode policies in the ADSL lines: modems could enter a *sleep mode* state with reduced energy consumption when no or little traffic is present.

Their implementation will not obviously be effective for all users. Indeed, users using P2P or watching IPTV service, continuously generate traffic, thus preventing the modem to enter the sleep mode.

Implementing sleep mode policies is important since different studies have shown that modems are consuming the same amount of energy independently from the offered traffic [16, 17]. ADSL2 standard [18] already incorporates recommendations about sleep mode policies, but at the moment no commercial modems implement them.

Other studies have already proposed some solutions to increase the energy efficiency in the ADSL access devices. In [3], authors propose a “roaming” scheme that allows end-users’ traffic to be carried by nearby Wi-Fi channels when the ADSL/Wi-Fi link of the user enters into sleep mode. In [19], energy savings obtained from the implementation of ADSL2 standards are evaluated. In this case, no coordination among modems is required and each physical link is managed separately. Unfortunately, the study is limited to a single test case.

In this work, we consider three power states: in the first one, both the DSLAM and users’ modems are completely turned

OFF, i.e., when entering into sleep mode the physical channel is lost; in the second one, the modems are able to enter a “low rate / low power” state in which the physical link is working at a low bit rate in the order of few kb/s with a reduced energy consumption; in the third one the modems are fully operative and they consume the maximum amount of power.

5.2.1. Simple Sleep Mode (SSM) policy

We start considering a simple policy that we call Simple Sleep Mode (SSM) which exploits a single line rate threshold, namely *MinRate*, to identify when the modem can enter into sleep mode.

Fig. 11(a) depicts the strategy reporting an example of line rate evolution over time. The modem monitors the line activity measuring the line rate. If the rate does not exceed *MinRate* for *Idle* seconds, then the line is assumed to be inactive and the modem enters into sleep mode. As soon as the line rate is above *MinRate*, the modem has to be woken up and to become fully operative. We consider that the wake up of the modem is almost instantaneous as envisaged by the ADSL2 standards.

5.2.2. Double threshold Sleep Mode (2SM) policy

The SSM policy is very conservative. When the modem is woken up for some traffic events, the modem has to wait for *Idle* seconds before enter into sleep mode again. In presence of background traffic, the line would remain active for long time even if little traffic is exchanged.

Therefore, it may be better that the modem, after being woken up to send background traffic, re-enters quickly into sleep mode state. The Double threshold Sleep Mode (2SM) policy exploits two traffic thresholds to distinguish between users’ activity and background traffic.

Fig. 11(b) shows an example of line rate evolution controlled by the 2SM policy. Similarly to SSM, as soon as the rate is higher than *Rate-low*, the modem is turned ON. Another threshold, *Rate-high*, is used to decide for how long the modem has to wait before entering into sleep mode. If the line rate is above *Rate-high* for at least T_{active} seconds, then the user is assumed to be active and the sleep mode can be entered only after an *Idle-long* inactivity period (first activity cycle in 11(b)). This means that the modem is left ON for *Idle-long* seconds even if the rate goes below *Rate-high* in a second moment. Otherwise, we assume that the traffic is likely to be related to background

Figure 11: Example of evolution of the line rate and the modem power consumption obtained using two sleep mode policies.

traffic and a fast sleep mode idle time, $Idle-short \leq Idle-long$, is used to quickly enter into sleep mode again.

5.2.3. Implementation issues of sleep mode policies at ADSL modems and impact on QoS

Several issues arise concerning the impact of the actual design and implementation of energy saving policies. In our work, we consider hypothetical energy saving policies, and do not enter into the details about their actual implementation.

For instance, any of the considered energy saving policies, at the high level, should be implemented without impacting on the Service-Level Agreements (SLAs) offered by the ISP. Alternatively, the ISP could also introduce a new SLA in which no energy saving policy is implemented, or it could incentivize the adoption of sleep mode enabled SLAs, sharing the benefit with its customers.

The implementation of sleep mode policies imply also some technical details related to the modems functioning. For example, re-synchronization delays are experienced each time the line has to be woken up from the sleep mode. These delays should be as short as possible, possibly imperceptible to the users. The actual impact of these delay is strictly related to the hardware implementation in the ADSL modems of the sleep mode policies. As already stated before, characterizing the implementation details of the sleep mode policies is out the scope of this work. Those are strictly related to the modem hardware and technological constraints. We just assume that the modems have enough buffering resources to compensate the variations of the line rate introduced by the sleep mode policies.

There are also other technical issues that need to be addressed. For instance, the detection of traffic during the sleep mode, and the transmission of background traffic have to be carefully considered.

The detection of traffic during sleep mode would depend on

the actual implementation. Some signaling channel must be available, e.g., using a low power link based on low power technologies or a wireless bluetooth sensor whose power consumption is negligible compared to the ADSL modem.

Alternatively, modems can offer a low power-low rate capability in the stand-by state, in such a way background traffic can be transmitted, e.g., software can check periodically for updates. Since the traffic transmitted for the check is usually limited, the modem would stay in low power state. Eventually, in case some important updates have to be installed, the large download would trigger the modem to full rate mode. Other possible solution is to transmit the traffic performing “offloading” over other available wireless networks [3].

Another issue concerns the interaction between the sleep mode and the session handling at the application layer. In general, the entering the stand-by state of the modem should not impact the session handling at the application layer. For instance, today, most web applications uses cookies as session indication, whose state is handled at the browser and it is not lost even in case of loss of connectivity. Other technique, like using keep-alive messages, can be used. If the system parameters (e.g., MinRate) of the sleep mode policies are correctly set, these keep-alive messages will allow to keep the ADSL line in low power/low bit-rate state without incurring into energy wastage.

6. Results

In this section we present the energy savings that could be achieved if previously described policies were in place.

6.1. Energy proportionality in the PoPs

In the following, the energy savings are computed considering the actual traffic profile which has been introduced in Sec. 2.

Figure 12: Energy savings for PoPs with linear proportional energy consumption.

Figure 13: Energy savings for PoPs with cubic proportional energy consumption.

For each PoP P , we suppose that an energy savings profile, $E(\rho)$, is in place.

6.1.1. Energy to traffic proportional technologies

The percentage of energy savings that can be achieved by using linear (Eq. 5) and cubic (Eq.6) energy proportional devices are reported, respectively, in Fig. 12 and in Fig. 13 for various values of the parameter α . When $\alpha = 0$, the consumption is not proportional at all. Therefore, no savings are possible and the consumption is the same as it is with current devices. As α increases, savings become larger: when consumption is fully load proportional ($\alpha = 1$, $E_{static} = 0$), up to 50-60% of energy can be saved in the linear case, while in the cubic case energy savings could reach percentages between the 80 and 90%. Although impressive, these results are theoretical and their validity have to be validated when energy proportional devices will be available, e.g., considering realistic values for α .

Since traffic patterns are very similar in all PoPs, savings are almost the same for the different PoPs.

6.1.2. Resource consolidation

The energy savings achieved by each PoP, when resource consolidation with 2-states is used, are shown in Fig. 14. In the plot, the threshold value m varies from 0 to 1.

Figure 14: Energy savings for PoPs with 2-states resource consolidation.

Figure 15: Energy savings for PoPs with 3-states resource consolidation.

When m is small (below 0.2 of the peak traffic) no savings are possible, since traffic rarely drops below m , the low capacity state is thus rarely entered. When $m > 0.8$, only the low capacity state is used, and savings depend on the actual value of m only. Depending on the traffic characteristics of the PoP, the optimal value of m slightly changes, but it is typically between 0.65 and 0.75. As shown in [20], the optimal value depends on the amount of energy that is saved in low capacity state and on the duration of the periods in which traffic is low and low capacity state can be entered. In summary, 20-35% of savings can be granted when 2-states resource consolidation schemes are in place.

The case of 3-states resource consolidation is shown in Fig. 15 varying m and fixing $l = m/2$. As expected, savings are higher than the 2-states case, yet, savings reach only 35-45%, i.e., 10-15% more than the 2-states resource consolidation case. Fig. 16 reports the savings for any value of l and m , with $l \leq m \leq 1$, for PoP B only. The optimal values of m and l are also in this case different for each PoP. We evaluated, through an exhaustive search, that values close to $m = 0.7$ and $l = 1/2m$ would guarantee good performance for all the PoPs (see the contour lines at bottom of Fig. 16).

Figure 16: Percentage of energy savings for PoP B with 3-states resource consolidation and different values of m and l .

Figure 17: CDF all T_{off} periods for all the lines. Results refer to SSM with $MinRate = 1$ kb/s.

6.2. Sleep mode policies at the ADSL lines

The evaluation of the effectiveness of the sleep mode policies in the ADSL lines is based on the following metrics:

- $T_{off}(i, l)$, the duration of the i -th sleep mode interval for the line l ;
- $T_{off-total}(l)$, the total time spent in the sleep mode by the line l corresponding to $T_{off-total}(l) = \sum_i T_{off}(i, l)$;
- $N_{trans}(l)$, the number of transitions from the sleep mode to the ON state occurred for a line l ;

Fig. 17 shows the CDF of all T_{off} time periods for all the lines considering SSM with $MinRate = 1$ kb/s. Different curves refer to different values of $Idle$. It is possible to notice that most of the periods last for just few hours. More in details, 80% of the periods in which the line is not active is below 6 hours. Note that shortening $Idle$ increases the number of short T_{off} periods, since, for smaller $Idle$, it is possible to enter sleep mode state more frequently, but the average duration of the T_{off} period is smaller. Varying the $MinRate$ we obtain similar CDFs.

The CDF of $T_{off-total}$ for SSM is depicted in Fig. 18. Also in this case we show the results for $MinRate$ equal to 1 kb/s. We

Figure 18: CDF of $T_{off-total}$ for all the lines. Results refer to SSM with $MinRate = 1$ kb/s.

Figure 19: Comparison of the average duration of total silent time $T_{off-total}$ using the SSM policy with different settings.

can appreciate that half of the lines can stay in sleep mode state more than 16 hours per day, even considering a conservative $Idle$ equal to 30 minutes. Selecting more aggressive settings of $Idle$ leads to higher savings. For instance, for $Idle$ equal to 5 minutes, half of the users is staying in the sleep mode state more than 20 hours per day, i.e., for more than 80% of the time. These results show that the implementation of sleep mode strategies can be effective for most of the users.

In Fig. 19, we report the average time spent by users in the sleep mode state, that is

$$E[T_{off-total}] = \frac{\sum_{l=1}^N T_{off-total}(l)}{N} \quad (10)$$

where N is the total number of users of our data set.

The same results have been retrieved for smaller users' subsets which have been randomly selected. The results on smaller subsets correspond to the results retrieved considering the whole set. This means that the users' behavior is independent from each other and the fact that users belong to the same PoP has no impact on the results.

To see the effect of *background traffic*, we consider different values of $MinRate$. Recall that a modem can enter into sleep mode only if the line rate has not exceeded $MinRate$ for $Idle$ seconds. In case no $MinRate$ is applied (e.g., $MinRate = 0$ kb/s), the average time spent in the sleep mode state is very limited. This can be clearly seen in Fig. 19. As $Idle$ is higher than 15

Figure 20: Comparing the number of transitions N_{trans} from sleep to ON state using the SSM policy with different settings.

Figure 21: CDF of the number of transitions N_{trans} from sleep to ON state using the SSM policy with $MinRate = 1\text{ kb/s}$ and different settings of $Idle$.

minutes, $E[T_{off-total}]$ is almost zero, meaning that *background traffic* is generated periodically and idle time periods are not long enough to allow the modem to entering sleep mode.

Instead, for $MinRate$ greater than zero, even a low value such as 1 kb/s brings significant improvements. Indeed, in the worst case with $Idle$ equal to 30 minutes, we have $E[T_{off-total}]$ equal to 15 hours. Decreasing $Idle$ it is possible to gain up to 3 hours of sleep mode considering $MinRate$ equal to 1 kb/s. If we further increase $MinRate$, no significant advantages are introduced; indeed, comparing 1 kb/s and 10 kb/s thresholds, the latter achieves a $E[T_{off-total}]$ about 2 hours longer in the best case. Thus, the value of $MinRate$ is not strongly affecting the time spent in the sleep mode state, provided it is larger than zero. It is then essential setting $MinRate$ to a value different from zero to cut off the background traffic. Indeed, the presence of these almost periodic and low bit rate traffic events is critical because it heavily impairs the performance of the sleep mode policies.

6.2.1. Impact on latency

The number of transitions from sleep mode to ON state is an important metric to evaluate the QoS for users. In fact, each transition implies a delay needed to re-synchronize the physical line or to allow modems to reach the maximum rate. Thus, it is better to choose wisely the settings in order to limit the number

Figure 22: Comparison of the average duration of total silent time $T_{off-total}$ using the 2SM policy with $Rate-low = 1\text{ kb/s}$ and $T_{active} = 2\text{ min}$.

of these transitions. In Fig. 20 we report the average number of state transitions $E[N_{trans}]$ for different values of $MinRate$ and $Idle$ considering the SSM policy. In case no $MinRate$ is applied, the number of transitions is very large for low values of $Idle$. Indeed, the line continuously switches from one state to the other. Instead, the number of state transitions decreases considerably for $T_{idle} \geq 20$ minutes, but the modems do not enter sleep mode anymore. Considering values of $MinRate$ equal to 1 kb/s or higher, we can see that there are not significant differences. In particular, we can see that the number of transitions is decreasing until $Idle$ equal to 15 minutes, while for larger values of $Idle$ the number of transitions is almost stable around 5. Thus, choosing a $Idle$ of 15 minutes is the best trade-off between the number of transitions and the time spent in the sleep mode state for the SSM policy.

In more details, Fig. 21 shows the CDF across lines of N_{trans} for $MinRate$ equal to 1 kb/s. Notice that most of the users experience a limited number of transitions. Indeed, considering T_{idle} equal to 15 minutes, 80% of users encounters 10 or less transitions during the day, while the maximum number of transitions experienced by an user is 60. This shows that few transitions in general would be imposed, provided that parameters are chosen wisely. Indeed, decreasing T_{idle} , the number of transitions experienced by the 80% of users increases to 25 transitions, while with T_{idle} set to 30 minutes, the same percentage of users has at most 7 transitions. Similar results are retrieved for higher values of $MinRate$ without any significant change.

The assumption that ADSL re-synchronization transient would be negligible is validated by the actual small number of transitions per day that an ADSL line has to face.

6.2.2. 2SM improvements

Considering the 2SM policy, the evaluation of $E[T_{off-total}]$ is shown in Fig. 22 for different settings. Given the previous considerations, we consider $Rate-low = 1\text{ kb/s}$, $Idle-long = 15\text{ min}$ and we set T_{active} equal to 2 min.

We notice that the values of $E[T_{off-total}]$ are almost constant with respect to $Rate-high$. Instead, decreasing $Idle-short$, it is possible to spend more time in the sleep mode state. In particular, using $Rate-high = 100\text{ kb/s}$ and $Idle-short = 2\text{ min}$, a line is silent for 18.2 h/day in average. For comparison, the SSM

Figure 23: Comparing the number of state transitions N_{trans} using the 2SM policy with $Rate-low = 1$ kb/s and $T_{active} = 2$ min.

Figure 24: Energy savings for the three models.

policy guarantees 16.3 h/day (highlighted in the plot for comparison). This means that with more aggressive settings it is possible to save more energy. However, these savings have to be compared with the number of state transitions needed that we report in Fig. 23. We can notice that 2SM setting introduces more state transitions with respect to the SSM case. This means that the 2 hours of additional savings implies additional impairments. More conservative settings could represent a better performance/gains compromise.

6.3. Achievable energy reduction and monetary savings

In the following, we compare the energy savings that can be achieved by the different strategies and we quantify which can be the energy saved in one-year long period for each PoP and how large can be the correspondent monetary savings.

6.3.1. Energy savings at the PoPs

The optimal performance of each proposed strategy for the PoPs is shown in Fig. 24. For the linear and cubic energy proportional technologies, we report the results for $\alpha = 1$ since this value guarantees the maximum possible savings. It can be considered as an upper bound of the energy savings that can be obtained with energy proportional schemes. For the resource consolidation schemes, instead, we perform an exhaustive search to find which are the best values of the thresholds in the case of the

Figure 25: Average energy consumption per users. SSM with $MinRate = 1$ kb/s and $T_{idle} = 15$ min. 2SM with $Rate-low = 1$ kb/s, $Rate-high = 100$ kb/s, $Idle-short = 5$ min, $Idle-long = 15$ min and $T_{active} = 2$ min.

2- and 3-states. Optimal values are different for each PoP. In the figure, the PoPs are ordered by increasing number of networking devices; notice that the size of the PoPs has little influence on energy savings. As expected, cubic proportional technology achieves the highest savings with more than 80% of energy that can be saved for any PoP. The savings obtained with linear proportionality are still consistent, typically about or higher than 50%. Even if a lot of effort is spent to develop network devices with energy proportionality characteristics, it will require some time before that they will become available on commercial devices. Instead, resource consolidation strategies are likely to be deployable in the near future and, even if lower than proportional schemes, the savings can be interesting, 20-30% of the 2-states case to around 40% with 3-states case. These solutions are probably very promising as short-term relief to high ISPs electricity bill. Note that between 2- and 3-states the energy savings present a small difference, thus it may be sufficient to implement 2-states resource consolidation to achieve interesting savings.

6.3.2. Energy savings at the ADSL lines

We evaluate the energy consumption for the SSM and 2SM policies and we compare them with respect to the case in which the lines are always active. We assume that the power consumption of a DSLAM modem port is around 1 W while an user's modem consumes about 5 W according to the values reported in [3, 16].

We consider two possible cases of the sleep mode state: (i) the modem is switched OFF, thus consuming approximately 0 W, (ii) the modem can vary the line rate reducing it to save energy. We select the following settings for the policies which represent the best trade-off between time spent in the sleep mode state and number of transitions:

SSM	$MinRate = 1$ kb/s, $Idle = 15$ min
2SM	$Rate-low = 1$ kb/s, $Rate-high = 100$ kb/s, $Idle-short = 5$ min, $Idle-long = 15$ min, $T_{active} = 2$ min

Using these settings we obtained the energy consumption reported in Fig. 25. The x-axis reports the ratio between the energy consumed by the modem in the sleep mode state and the

energy consumed in the ON state. Thus, $x=0$ is the case in which the modem is powered OFF when it enters sleep mode. When this ratio is equal to 1 it means that the system does not support sleep mode.

The figure reports for comparison the amount of energy consumed by modems for the SSM and for the 2SM policies. The energy is evaluated with two different scales. Left y-axis reports the average amount of energy consumed per user in one day considering the average values obtained from the analysis of the dataset. Instead, right y-axis reports the energy consumed in one year assuming 1 million of users. We can notice that the energy consumed by the two strategies is very similar, with 2SM performing a little better.

If we refer to the “Code of Conduct (CoC) on Energy Consumption of Broadband Equipment” to get realistic power consumption values of user modem in the ON and sleep mode states¹¹, we have that the ratio is about 0.54. In this case, SSM allows to easily save about 32% of energy, which consists to save around 16.5 GWh in one year for a population of 1 million of users. Note that, since the DSLAM modem accounts for 1/6 of the energy consumption of a line, an operator that is willing to implement a sleep mode strategy only at DSLAMs can save about 2.75 GWh, which corresponds to 310 kEuro per year, assuming a price of electricity equal to 0.112 Euro/kWh¹². If we consider the best scenario (i.e., the modem does not consume during the sleep mode), the energy saved is about 38 GWh, guaranteeing to double the savings.

6.3.3. Total energy and monetary savings per PoP in one-year long period

In order to get as much as possible realistic savings, we choose the most conservative policies. For what concerns the PoPs, we estimate the savings in the case of the 2-states resource consolidation strategy. The value of threshold m has been set equal to the optimal one, as in Fig. 24. We assume that the energy consumed in the High state corresponds to the current consumption of networking devices from Table 2. For the ADSL lines, we select the SSM policy and we use the same setting of Fig. 25, thus $MinRate = 1$ kb/s and $T_{idle} = 15$ min.

Table 3 reports, respectively, the energy savings at each PoP, the energy savings in the DSLAMs associated to a PoP, the total energy savings that can be achieved in a PoP and the correspondent total monetary savings. The price of electricity is assumed 0.112 Euro/kWh.

Note that the amount of saved energy, and of money, is substantial in the larger PoPs. This means that only in these PoPs there is an economical convenience to invest in the implementation of energy efficient policies.

Looking at the energy saved by applying the SSM policy at the DSLAMs modems, it possible to notice, in the PoPs with a large number of ADSL users, that the saved energy is significant even if the DSLAMs modems consume only about 1 W.

ID	S_{PoP} [MWh]	S_{ADSL} [MWh]	S_{ToT} [MWh]	Money savings [kEuro/year]
A	8.5	-	8.5	0.95
B	153.8	103.8	257.7	28.96
C	85.9	67.8	153.8	17.2
D	60	7	67.1	7.5
E	125	3.3	128.3	14.4
F	352.3	165.6	517.87	58
G	195.4	174.2	369.64	41.4
H	103.7	67.1	170.84	19.1

Table 3: Energy savings at the PoPs

7. Conclusions

We exploited a large dataset of real measurements, collected in the operative network of FASTWEB, a national-wide Italian ISP, to perform an extensive analysis of the energy and traffic of the PoPs and to characterize the traffic profiles of the ADSL lines. Next, we evaluated the achievable savings considering energy efficient policies and actual collected data.

Our analysis confirmed that the energy consumption of the PoPs is independent from traffic load while is strongly correlated to the external temperature. We thus considered to implement at the PoPs solutions based on energy-to-traffic proportional technologies or on the resource consolidation practice. These solutions can adapt the energy consumption of the PoPs to their current traffic load. Results also show that ADSL lines are typically inactive for long time periods. We thus evaluated sleep mode policies that, according to the line traffic activity, can decide to put the modem into a sleep mode state characterized by a reduced power consumption and transmitting capabilities.

We estimated that the implementation of these policies can effectively improve the energy efficiency. It is possible to save about 30% of energy at the PoPs and in the access network using conservative settings. The obtained results show also that interesting monetary savings can be achieved thanks to the reduction of the ISPs’ energy bill and these savings can motivate the implementation of these policies and partially cover the required investments.

Acknowledgment

The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n. 257740 (Network of Excellence “TREND”).

- [1] C. Lange, D. Kosiankowski, R. Weidmann, A. Gladisch, Energy consumption of telecommunication networks and related improvement options, Selected Topics in Quantum Electronics, IEEE Journal of 17 (2011) 285–295.
- [2] J. Baliga, R. Ayre, K. Hinton, W. V. Sorin, R. S. Tucker, Energy Consumption in Optical IP Networks, J. Lightwave Technol. 27 (2009) 2391–2403.
- [3] E. Goma, M. Canini, A. Lopez Toledo, N. Laoutaris, D. Kostić, P. Rodriguez, R. Stanojević, P. Yagüe Valentin, Insomnia in the access: or how to curb access network related energy consumption, SIGCOMM Comput. Commun. Rev. 41 (2011) 338–349.

¹¹<http://iet.jrc.ec.europa.eu/energyefficiency/ict-codes-conduct/energy-consumption-broadband-communication-equipment>

¹²<http://epp.eurostat.ec.europa.eu>

- [4] E. Bonetto, M. Mellia, M. Meo, Energy Profiling of ISP Points of Presence, in: Communications (ICC), 2012 IEEE International Conference on, Ottawa, Canada, pp. 5973–5977.
- [5] E. Bonetto, A. Finamore, M. Munafo, R. Fiandra, Sleep mode at the edge: How much room is there?, in: Telecommunications Network Strategy and Planning Symposium (NETWORKS), 2012 XVth International, Rome, Italy, pp. 1–6.
- [6] A. Finamore, M. Mellia, M. Meo, M. M. Munafò, D. Rossi, Experiences of Internet Traffic Monitoring with Tstat, IEEE Network 25 (2011) 8–14.
- [7] J. L. Rodgers, W. A. Nicewander, Thirteen ways to look at the correlation coefficient, The American Statistician 42 (1988) 59–66.
- [8] C. Malone, C. Belady, Metrics to characterize data center & it equipment energy use, in: Proceedings of the Digital Power Forum, Richardson, TX, USA.
- [9] C. Belady, A. Rawson, J. Pfeuger, T. Cader, Green Grid Data Center Power Efficiency Metrics: PUE and DCIE, White Paper 6, The Green Grid, 2007.
- [10] 2012 Uptime Institute Data Center Industry Survey, 2012.
- [11] J. Chabarek, J. Sommers, P. Barford, C. Egan, D. Tsang, S. Wright, Power awareness in network design and routing, in: IEEE INFOCOM 2008, Phoenix, USA.
- [12] D. Abts, M. R. Marty, P. M. Wells, P. Klausler, H. Liu, Energy proportional datacenter networks, SIGARCH Comput. Archit. News 38 (2010) 338–347.
- [13] R. Bolla, R. Bruschi, A. Cianfrani, M. Listanti, Enabling backbone networks to sleep, Network, IEEE 25 (2011) 26–31.
- [14] L. Chiaraviglio, M. Mellia, F. Neri, Minimizing ISP Network Energy Cost: Formulation and Solutions, IEEE/ACM Transactions on Networking (2011).
- [15] K. Choi, Dynamic voltage and frequency scaling for energy-efficient system design, Ph.D. thesis, Los Angeles, CA, USA, 2005. AAI3180330.
- [16] A. O. Calchand, P. Branch, J. But, Analysis of Power Consumption in Consumer ADSL Modems, Technical Report 100125A, Centre for Advanced Internet Architectures, Swinburne University of Technology, Melbourne, Australia, 2010.
- [17] A. Valenti, F. Matera, G. Belevi, Power consumption measurements of access networks in a wide geographical area test bed and economic perspectives, in: Future Network Mobile Summit (FutureNetw), 2012.
- [18] G.992.3: Asymmetric Digital Subscriber Line transceivers 2 (ADSL2), 2009.
- [19] A. Gladisch, C. Lange, R. Leppla, Power Efficiency of Optical versus Electronic Access Networks, in: Optical Communication, 2008. ECOC 2008. 34th European Conference on, pp. 1–4.
- [20] M. A. Marsan, L. Chiaraviglio, D. Ciullo, M. Meo, Optimal Energy Savings in Cellular Access Networks, in: Proceedings of the 1st International Workshop on Green Communications (GreenComm '09), Dresden, Germany.